arXiv:1410.7103v2 [cs.LO] 6 Nov 2014

Confusion in the Church-Turing Thesis

Barry Jay and Jose Vergara
University of Technology, Sydney
{Barry.Jay,Jose.Vergara}@uts.edu.au

November 7, 2014

Abstract

The Church-Turing Thesis confuses numerical computations with symbolic computations. In particular, any model of computability in which equality is not definable, such as the λ-models underpinning higher-order programming languages, is *not* equivalent to the Turing model. However, a modern combinatory calculus, the SF-calculus, can define equality of its closed normal forms, and so yields a model of computability that is equivalent to the Turing model. This has profound implications for programming language design.

DRAFT

1 Introduction

The λ-calculus [17, 4] does not define the limit of expressive power for higher-order programming languages, even when they are implemented as Turing machines [69]. That this seems to be incompatible with the Church-Turing Thesis is due to confusion over what the thesis is, and confusion within the thesis itself. According to Soare [63, page 11], the Church-Turing Thesis is first mentioned in Steven Kleene's book *Introduction to Metamathematics* [42]. However, the thesis is never actually stated there, so that each later writer has had to find their own definition. The main confusions within the thesis can be exposed by using the book to supply the premises for the argument in Figure 1 overleaf. The conclusion asserts the λ-definability of equality of closed λ-terms in normal form, i.e. syntactic equality of their deBruijn representations [20]. Since the conclusion is false [4, page 519] there must be a fault in the argument.

The basic error is introduced by statement (2), since Turing's proof was for *numerical* functions, i.e. functions acting on natural numbers or positive integers, and not for *symbolic* functions, i.e. functions acting on words in some alphabet, usually well-formed formulas, such as λ-terms. So, there is a confusion between numerical interpretation and symbolic interpretation. Since the link to higher-order programming concerns symbolic functions, we should replace (2) by a less ambiguous statement, such as:

(2′) The computable symbolic functions are equivalent to the λ-definable symbolic functions.

Now the focus shifts to the nature of the equivalence in (2′) and its use when inferring (3). We will see that the traditional arguments do indeed induce a

(1)	"... every function which would naturally be regarded as computable is computable under his [Turing's] definition, i.e. by one of his machines ..."	(page 376)
(2)	"The equivalence of the computable to the λ-definable functions ... was proved by Turing 1937."	(page 320)
(3)	Every function which would naturally be regarded as computable is λ-definable.	(equivalence)
(4)	The equality of λ-terms in closed normal form is λ-definable.	(specialise)

Figure 1: A faulty argument, with premises quoted from Kleene [42]

relationship between a λ-model of computability and a Turing model, in that each can simulate the other, but this mutual simulation is too weak to be an equivalence, to support the substitution of "λ-definable" for "computable" in the argument.

Thus we have three confusions. The faulty argument exposes the confusion of numerical functions with symbolic ones, and of mutual simulation with equivalence. Then there is the confusion about what, exactly, the Church-Turing Thesis is.

We will explore these confusions in three stages. First, we will explore the confusion surrounding Church's Thesis, as introduced by Kleene [42]. This will lead us to formalise models of computability and their relative expressive power. It will follow that any λ-model of computability that is suitable for modeling higher-order programming languages is less expressive than the recursive model. Second, we will extend the analysis to Turing's Thesis and the Church-Turing Thesis, and their impact on programming language design. Third, we will show that a recent combinatory calculus, the SF-calculus [34] yields a model of computability that is both suitable for higher-order programming and equivalent to the Turing model. By supporting a form of intensional computation, it suggests a more powerful approach to programming language design.

The sections of the paper are: Section 1 Introduction; Section 2 Church's Thesis; Section 3 Models of computability; Section 4 Comparison of models; Section 5 Turing's Thesis; Section 6 The Church-Turing Thesis; Section 7 Programming language design; Section 8 Intensional computation; and Section 9 Conclusions.

2 Church's Thesis

The background to Church's Thesis is Church's paper of 1936 *An unsolvable problem of number theory* [16] which, alongside Alan Turing's work discussed later, showed that some numerical problems do not have a computable solution, so that Hilbert's decision problem does not have a general solution. It is clear that Church's focus was on numerical functions, as all his formal definitions were expressed in numerical terms. For example, he writes, "A function F of one positive integer is said to be λ-definable if ..." [16, page 349]. Again, in Theorem XVII he writes "Every λ-definable function of positive integers is recursive".

DRAFT

That said, he does broaden the discussion when considering *effective calculability*. On the one hand, an *effectively calculable* function of positive integers is defined to be "a recursive functions of positive integers" or "a λ-definable function of positive integers". On the other hand, he writes:

> ... [in] some particular system of symbolic logic ... it is necessary that each rule of procedure be an effectively calculable operation, ... Suppose we interpret this to mean that, in terms of a system of Gödel representations for the expressions of the logic, each rule of procedure must be a recursive operation, ...

That is to say, a symbolic function is recursive if it can be simulated by a numerical function that is recursive, where simulation is defined using Gödel numbering. Further, the difference between the numerical and symbolic domains does not seem to be important to him, as he writes, in a footnote:

> ... in view of the Gödel representation and the ideas associated with it, symbolic logic can now be regarded, mathematically, as a branch of elementary number theory.

In contrast to Church's narrow constructions, Kleene's definitions have a broad scope. For example, he gives three definitions of λ-definability in his paper *λ-definability and recursiveness* [41], according to whether the domain of definition is the non-negative integers, the λ-terms themselves, or other mathematical entities for which a λ-representation has been given. It is the interplay between these definitions that is the primary cause of confusion.

Kleene introduces Church's Thesis as Thesis I in his 1952 book (Section 60, page 300) as follows:

(CT) Every effectively calculable function (effectively decidable predicate) is general recursive.

The crucial question is to determine the domain of the effectively calculable functions.

At the point where Church's Thesis is stated, in Section 60, the general recursive functions are all numerical, so it would seem that the effectively calculable functions must also be numerical. However, he does not include the phrase "of positive integers" in his statement of the thesis, in the careful manner of Church. We are required to add this rider in order to make sense of the thesis.

Later, in Section 62. Church's Thesis, Kleene presents seven pages of arguments for the thesis, which he groups under four headings A–D. In "(B) Equivalence of diverse formulations" he asserts that the set of λ-definable functions and the set of Turing computable functions are "co-extensive" with the set of general recursive functions. Again, this only makes sense if the functions are presumed to be numerical functions. This paragraph is also the source of statement (2) from Figure 1.

If we were to stop at this point, then the explanation of the faulty argument would be quite simple: statement (2) should be read in a context where all functions are numerical, or be replaced by "The equivalence of the computable to the λ-definable **numerical** functions was proved by Turing 1937". The restriction to numerical functions propagates to statement (3) which cannot then be specialised to equality of λ-terms.

DRAFT

However, in "(D) Symbolic logics and symbolic algorithms", Kleene reprises Church's definition of symbolic functions that are recursive, so that, by the end of Section 62, we have two definitions of effectively calculable functions and of recursive functions. So there are two versions of Church's Thesis, one for numerical functions (NCT) and one for symbolic functions (SCT). Unpacking the definition of a general recursive symbolic function to make the role of simulations explicit, the two versions become:

(NCT) Every effectively calculable numerical function (effectively decidable numerical predicate) is general recursive.

(SCT) Every effectively calculable function (effectively decidable predicate) can be simulated by a function which is general recursive.

Summarising, we see that Church was careful to separate theorems about numerical functions from the discussion of computation in symbolic logic. By contrast, Kleene presents a single statement of Church's Thesis with evidence that confuses the numerical with the symbolic. In turn, this confuses two different questions: whether two sets of numerical functions are the same; and whether there is an encoding that allows functions in a symbolic logic to be simulated by recursive functions on numbers. These confusions can be defused by isolating two versions of Church's Thesis which, from the viewpoint of Post [57], qualify as scientific laws. Now this distinction is enough to eliminate the numerical version of the faulty argument but the symbolic version remains. To eliminate it, we must show that an equivalence of models of computation requires more than their mutual simulation. To make this precise, we must formally define models of computability, their simulations and equivalence.

3 Models of computability

We adopt the simplest definition of model of computability in which the discussion of simulation makes sense. This was introduced by Boker and Dershowitz as a *model of computation* [9], but since the focus is upon functions for which a computation is possible, rather than the actual mechanics of computation, it seems more accurate to call them models of computability. Note, too, that the domain of the computable functions is not actually required to be symbolic in any way, or even to be enumerable, however natural this may be. This makes the notion too weak for many purposes, but here it emphasises the generality of our results.

A *model of computability* $(D, \mathcal{F})$ is given by a domain D which is a set of *values* that provides arguments and results of computations, and a collection $\mathcal{F}$ of partial functions from powers of D to D. Here are some examples.

The partial recursive functions on natural numbers form a model with domain given by the natural numbers and functions given by the partial recursive functions. Call this the *recursive model of computability*.

Recall that an injective function is a total function that does not identify distinct arguments. For any finite alphabet Σ, and any domain D equipped with an injective function from D to the words of Σ, the *Turing model of computability on* D has domain D and partial functions given by those which can be computed by a Turing machine with alphabet Σ. When the choice of domain is understood

DRAFT

from the context, or unimportant, then it may be called the *Turing model of computability*.

Any *applicative rewriting system* [67] has a *normal model* whose values are the closed terms in normal form, and whose partial functions are those representable by closed terms in normal form. Further, any subset D of values determines a model, where the partial functions are now defined on a value in D only if the result is also in D.

For example, classical combinatory logic has terms built from applications of the operators S and K and variables. As well as its normal model, there is a *numerical model* whose domain is restricted to be the Church numerals. Also, one can use Polish notation to encode combinators as words using the alphabet $\Sigma = \{A, S, K\}$, where A is for application. For example, $S(KK)$ is mapped to $ASAKK$. This yields a Turing model of computability for SK-normal forms.

Similarly, λ-calculus has *normal models* and *numerical models*, once the terms and reduction rules have been specified. First, a λ-term is unchanged by renaming of bound variables, i.e. is an α-equivalence class in the syntax [4]. Since this equivalence is a distraction, we will work with λ-terms using deBruijn notation [20] so that, for example, $\lambda x.x$ becomes $\lambda 0$ and $\lambda x.\lambda y.xy$ becomes $\lambda\lambda 10$. Second, there are various choices of reduction rules possible, with each choice producing a normal λ-model. Define a *λ-model of higher-order programming* to be any normal λ-model in which equality of closed normal forms is not definable. This excludes any model whose domain is numerical, and would seem to include any models that are relevant to higher-order programming. Certainly the λ-models of higher-order programming include those given by β-reduction alone, or $\beta\eta$-reduction.

4 Comparison of models

Now consider what it means for one model of computability to be more expressive than another. The simple interpretation requires that their computable functions have the same domain and then compares sets of computable functions by subset inclusion, as is done by John Mitchell [49] and Neil Jones [39]. The choice of domain is important here. For example, if the domain consists of natural numbers then the λ-model and the recursive model are indeed equivalent. However, this restriction is unreasonable for modeling higher-order programming since functions must be among the values. Now it is an easy matter to see that any normal λ-model of computability has fewer computable functions than the Turing model. For example, no λ-term can decide equality of values but this function is in the Turing model.

The complex interpretation of relative expressive power allows the domains to vary, but now the comparison of computability over domains D_1 and D_2 must be mediated by simulations, which are given by encodings. Note that it makes no sense to consider a simulation in one direction, and compare sets of functions in the other direction, since, as observed by Boker and Dershowitz [9], this can lead to paradoxes. Rather, there should be encodings of each domain in the other that are, in some sense, inverse to each other. Various choices are possible here but two requirements seem to be essential. First, the encodings should be injective functions, since distinct values should not be identified. Second, the encodings should be *passive*, in the sense that they are not adding expressive

DRAFT

power beyond that of their target model. In particular, they should be effectively calculable. This requirement can only be verified informally, on a case by case basis, since functions from D_1 to D_2 are not in the scope of either model.

A *simulation* of a model of computability $(D_1, \mathcal{F}_1)$ in another such $(D_2, \mathcal{F}_2)$ is given by an injective *encoding* $\rho : D_1 \to D_2$ such that every function f_1 in $\mathcal{F}_1$ can be *simulated by* a function f_2 in $\mathcal{F}_2$ in the following sense: for all $x_1, \ldots, x_n \in D_1$ such that $f_1(x_1, \ldots, x_n)$ is defined, then $f_2(\rho(x_1), \ldots, \rho(x_n))$ is defined and

$$\rho(f_1(x_1, \ldots, x_n)) = f_2(\rho(x_1), \ldots, \rho(x_n)) \ .$$

For example, Gödelisation provides a simulation of the normal SK-model into the recursive model. More generally, Church and Kleene both use this notion of simulation to define recursive symbolic functions. Gödelisation seems to be passive.

Further, the encoding of the natural numbers using Church numerals provides a simulation of the recursive model into the normal SK-model or any normal λ-model.

Other related notions of simulation can be found in the literature, e.g. [39, 9]. For example, Richard Montague [52, page 430] considers, and Hartley Rogers [59, page 28] adopts, a slightly different approach, in which the encoding of numbers is achieved by reversing a bijective Gödelisation. However, this inverse encoding may not be a simulation. For example, the equality of numbers cannot be simulated by a λ-term over the domain of closed λ-terms in normal form. Rogers, like Kleene, ensures a simulation by *defining* the computable functions in the symbolic domain to be all simulations of partial recursive functions, but this says nothing about λ-definability.

Given the formal definition of simulations, it may appear that the strongest possible notion of equivalence is that each model simulates the other. However, if the encodings are passive then so are the *recodings* from D_1 to D_1 and from D_2 to D_2 obtained by encoding twice. Since these are in the scope of the two models, we can require that recodings be computable.

Let $(D_1, \mathcal{F}_1)$ and $(D_2, \mathcal{F}_2)$ be two models of computability with simulations $\rho_2 : D_1 \to D_2$ and $\rho_1 : D_2 \to D_1$. Then $(D_2, \mathcal{F}_2)$ is *at least as expressive* as $(D_1, \mathcal{F}_1)$ if the recoding $\rho_2 \circ \rho_1 : D_2 \to D_2$ is computable in $\mathcal{F}_2$. If, in addition, $(D_1, \mathcal{F}_1)$ is more expressive than $(D_2, \mathcal{F}_2)$ then the two models are *weakly equivalent*. Note that this notion of equivalence is indeed an equivalence relation on models of computability.

It is interesting to compare this definition with those for equivalence of *partial combinatory algebras* by Cockett and Hofstra [19] and John Longley [45]. They would not require the encodings to be injective, but Longley would require that the recodings be invertible. Adding the latter requirement is perfectly reasonable but is immaterial in the current setting.

It is easy to prove that the recursive model is at least as expressive as any λ-model. Our focus will be on the converse.

Theorem 1 *Any model of computability that is at least as expressive as the recursive function model can define equality of values.*

Proof The recursive model is presumed to use 0 and 1 for booleans. In the other model, identify the booleans with the encodings of 0 and 1 so that

DRAFT

the equality function is given by recoding its arguments and then applying the simulation of the equality of numbers. □

Corollary 2 *The normal model of computability for SK-calculus is not weakly equivalent to the recursive function model.*

Proof If the normal SK-model could define equality then it could distinguish the values SKK and SKS but the standard translation from combinatory logic to λ-calculus identifies them (both reduce to the identity), and so they cannot be distinguished by any SK-combinator. □

Corollary 3 *No λ-model of higher-order programming is weakly equivalent to the recursive model of computability.*

Proof Since normal λ-models do not define equality, the result is immediate. Note that Longley has proved the analogous result for his definition of equivalence [45]. □

5 Turing's Thesis

Turing's paper of 1936 *On Computable Numbers, with an application to the Entscheidungsproblem* was, like Church's paper, concerned with numerical computation and Hilbert's decision problem. Like Church, Turing was careful to limit his definitions, e.g. of *computable functions*, to numerical functions while showing awareness of a broader scope. For example, in the first paragraph he writes:

> Although the subject of this paper is ostensibly the computable *numbers*, it is almost equally easy to define and investigate computable functions of an integral variable, or a real or computable variable, computable predicates, and so forth.

Similarly, the use of an unspecified alphabet of symbols on the tape of a Turing machine encourages us to consider computation over arbitrary symbolic domains.

Once again, Kleene confuses these two meanings in his third piece of evidence for Church's Thesis, headed (C): Turing's concept of a computing machine [42, page 320] where Kleene writes "Turing's computable functions (1936-7) are those which can be computed by a machine of a kind which is designed, according to his analysis, to reproduce all the sorts of operations which a human computer could perform, working according to preassigned instructions." As we have seen above "Turing's computable functions" are, by definition, numerical, while a human computer faces no such restriction.

In more compressed form, this confusion re-appears in Kleene's statement of Turing's Thesis. It is given in Section 70. Turing's thesis, as a sub-ordinate clause of the opening statement which, when elevated to an independent thesis, becomes:

(TT) Every function which would naturally be regarded as computable is computable under Turing's definition, i.e. by one of his machines.

DRAFT

Now the phrase "every function which would naturally be regarded as computable" surely includes all computations in formal systems such as λ-calculus or combinatory logic. For example, it would be a distortion to assume that "naturally" here refers to the natural numbers. On the other hand, "Turing's definition" is certainly numerical. As with Church's Thesis, the solution is to create a numerical thesis (NTT) and a symbolic one (STT) as follows:

(NTT) Every numerical function which would naturally be regarded as computable is computable under Turing's definition.

(STT) Every function which would naturally be regarded as computable can be simulated by a function which is computable under Turing's definition.

From the viewpoint of Post [57], both versions of the thesis will qualify as scientific laws. Now we can express some classical results in the new terminology.

Theorem 4 *Turing's Numerical Thesis is logically equivalent to Church's Numerical Thesis.*

Proof Apply Kleene's 30th theorem, i.e. Theorem XXX [42, page 376]. □

Theorem 5 *The recursive model of computability is weakly equivalent to any Turing model of computability.*

Proof The traditional simulations yield encodings that are computable. □

Corollary 6 *Turing's Symbolic Thesis is logically equivalent to Church's Symbolic Thesis.*

Proof Any simulation into a Turing model yields a simulation into the recursive model by composing with the simulation given by weak equivalence. The converse is similar. □

Corollary 7 *No λ-model of higher-order programming is weakly equivalent to the Turing model.*

Proof Since weak equivalence is transitive, the result follows from Corollaries 3 and 5. □

Now any reasonable notion of equivalence must imply weak equivalence. So it follows that the Turing model of computability is strictly more expressive than any λ-model of computability suitable for modeling higher-order programming.

6 The Church-Turing Thesis

The mathematical confusions are now defused, with separate numerical and symbolic versions of Church's Thesis and of Turing's Thesis, and a clear account of equivalence of models. In turn, this must require two versions of the Church-Turing Thesis. Putting this defusion to one side, there remains some confusion about what, exactly, the Church-Turing Thesis is since, although introduced in Kleene's book [42, page 382], it is nowhere defined. We have evidence for four accounts.

Among the statements in Kleene's book, the closest candidate is the opening of Section 70:

DRAFT

> Turing's thesis that every function which would naturally be regarded as computable is computable under his definition, i.e. by one of his machines, is equivalent to Church's thesis by Theorem XXX.

On first reading, it is rather difficult to determine the nature of this declaration, as it contains the statement of Turing's thesis (TT) plus the statement of a theorem,

> Turing's Thesis is equivalent to Church's Thesis.

with its proof "by Theorem XXX". If this is the Church-Turing Thesis, then it is a theorem asserting logical equivalence of two theses. The other candidate statement is the numerical version of (2) from Figure 1 which is also a theorem, but this time asserting mathematical equivalence of two models.

According to Solomon Feferman [22], the Church-Turing Thesis was born in Alonzo Church's 1937 review of Alan Turing's paper on computability [69] which declared

> As a matter of fact, there is involved here the equivalence of three different notions: computability by a Turing machine, general recursiveness in the sense of Herbrand-Gödel-Kleene, and λ- definability in the sense of Kleene and the present reviewer.

If by "notion" is meant a model of computability, then Church's statement is about equivalence of models. Feferman goes on to say

> Thus was born what is now called the Church-Turing Thesis, according to which the effectively computable functions are exactly those computable by a Turing machine. The (Church-)Turing Thesis ...

Now Feferman identifies the Church-Turing Thesis with the (Church)-Turing Thesis with Kleene's account of Turing's Thesis. It seems that Feferman considers this to be a single thesis with two (or three) names.

Finally, the literature of the last fifty years has thrown up many versions of the theses, e.g. [25, 62, 5]. The best way to make sense of this variety is to view the Church-Turing Thesis as the class of all statements that are logically equivalent to Church's Thesis or to Turing's Thesis. In this manner, all of the theses and proofs of logical equivalence are gathered under a single heading. This broad interpretation may explain why Kleene did not give a statement of it. In any event, this broad interpretation seems most appropriate when considering the impact of the Church-Turing Thesis on programming language design.

7 Programming language design

Here are four examples, from the last fifty years, of how confusion in the Church-Turing Thesis has limited the design space for programming languages.

Peter Landin's seminal paper of 1966 *The Next 700 Programming Languages* [44] proposes a powerful model of programming language development in which λ-calculus is the universal intermediate language. That is: create a source language with various additional features such as types, or let-declarations; transform this into λ-calculus; then implement an evaluation strategy for λ-calculus

DRAFT

(e.g. lazy or eager) as a Turing machine. His main comment about the suitability of λ-calculus for this role is:

> A possible first step in the research program is 1700 doctoral theses called "A Correspondence between x and Church's λ-notation."

which is footnoted "A not inappropriate title [for this paper] would have been "Church without lambda."" So, Landin sees a central role for λ-calculus, with a research program that would occupy a generation of computer scientists. The simplest interpretation of "Church without lambda" is "Church's Thesis without lambda". This research program influenced many languages designs, including Algol68 [58], Scheme [65], and ML [48].

Matthias Felleisen, in his paper of 1990 on the expressive power of programming languages [23], comments:

> Comparing the set of computable functions that a language can represent is useless because the languages in question are usually universal; other measures do not exist.

After this appeal to the Church-Turing Thesis, the paper goes on to consider various λ-calculi, in the belief that nothing has been left out. Although Felleisen makes some useful distinctions, his paper excludes the possibility of going beyond λ-calculus, which limits its scope.

Henk Barendregt et al [6] present a current version of the Church-Turing Thesis:

> **Church-Turing Thesis** The notion of intuitive computability is exactly captured by λ-definability or by Turing computability.

It overstates the significance of λ-calculus, since "exactly captured" suggests an equivalence of models, that goes beyond the existence of mutual simulations. It is a paraphrase of statement (2) from the faulty argument.

Robert Harper's book *Practical Foundations for Programming Languages* [30] contains a version of Church's Thesis (given as Church's Law) which is careful to limit its scope to natural numbers. It also emphasises that equality is not λ-definable (given as Scott's Theorem). However, while the title proclaims the subject matter to be the foundation for programming languages in general, it is solely focused on the λ-calculus, with no allowance made for other possibilities. If there remains any doubt about the author's views about foundations, consider the following slogan, attributed to him by Dana Scott in the year of the book's publication. In a talk during the Turing Centenary celebrations of 2012, he asserts [61]:

λ conquers all!

There is no explicit justification given for this focus, which we can only assume is based upon Landin's research program, and the Church-Turing Thesis.

Here are some examples of calculi and languages that don't easily fit into Landin's program, since they may exceed the expressive power of λ-calculus. Candidates include: first-order languages such as SQL [15]; languages without an underlying calculus, such as Lisp [47] with its operators `car` and `cdr`; and the intensional programming language Rum [66]. Richer examples include the *self-calculus* of Abadi and Cardelli for object-orientation [1], the *pattern calculus*

DRAFT

for structure polymorphism [38], the *pure pattern calculus* for generic queries of data structures [35, 33]. Again, the **bondi** programming language [10] uses pure pattern calculus to support both generic forms of the usual database queries, and a pattern-matching account of object-orientation, including method specialisation, sub-typing, etc. Most recently, *SF-calculus* is a combinatory calculus that extends generic queries from data structures to functions of all kinds [34]. The approach supports definable equality of closed normal forms, typed self-interpreters [36] (see also [56]), and Guy Steele's approach [64] to growing a language [37]. It can also be extended to a concurrent setting [27, 26]. The richer calculi above have not been shown equivalent to λ-calculus. Rather, all evidence points the other way, since the factorisation operator F of SF-calculus cannot be defined in SK-calculus [34].

Summarising, while the ability to simulate λ-calculus as a Turing machine has been enormously fruitful, the larger claims of the Church-Turing Thesis have been suggesting unnecessary limits on programming language design for almost fifty years.

8 Intensional computation

Having exposed a gap between the expressive power of λ-calculus and of Turing machines, it is natural to consider how to bridge it.

It is not a simple matter to overcome the limitations of λ-calculus by, say, adding an operator for equality. The essential difficulty is that while λ-terms describe algorithms, i.e. capture *intensions*, this intensional information cannot be recovered from within λ-calculus in any uniform manner. Rather λ-terms can only extract *extensional* information about input-output behaviour [17, page 2]. To redress this, various efforts have been made, in the context of partial evaluation and decompilation, to extend λ-calculus with Gödelisation. This has been done for simply typed λ-calculus [7], a combinatory calculus [29], and untyped λ-calculus augmented with some labels [51]. However, some of the attractive features of pure λ-calculus, such as being typable, confluent, and a rewriting system have been compromised.

Alternatively, the development of intensional computation can begin afresh. Intensionality has been the subject of much research by philosophers [52, 21, 14, 40], logicians [24, 68, 12], type theorists [46, 53, 11] and computer scientists [31, 3, 13], so before proceeding, let us determine what it will mean for us. In the concrete setting of λ-calculus, when should two λ-terms be considered intensionally equal? Should this be limited to closed normal forms or are arbitrary terms to be included? In part, the answer depends upon whether your semantics is denotational or operational.

Denotational semantics constructs the meaning of a program from that of its fragments, whose contexts may supply values to free variables, or determine whether or not the evaluation of the fragment terminates. Examples may be found in *domain theory* [43, 60], *abstract and algebraic data types* [28, 50], the *effective topos* [32], and *partial combinatory algebras* [8, 18, 45]. This suggests that all terms are included but equality of arbitrary lambda terms is not computable [4, page 519].

By contrast, other semantics do not account for terms without normal form or for open terms, and so avoid the need to assign them values. For exam-

DRAFT

ple, *axiomatic recursion theory* [59, 39] uses Kleene equality [42, page 327]. Again, *operational semantics* in the style of Gordon Plotkin's *structured operational semantics* [55] can limit its values to be closed terms that are, in some sense, normal, e.g. are irreducible, or in head-normal form [4], etc. Thereby, various problems caused by non-termination, such as the difficulty of defining the parallel-or function [2], do not arise. In particular, it is easy to represent equality.

Thus, the challenge is to extend standard calculi with the ability to query the internal structure of closed normal forms, in a process akin to Gödelisation, while retaining many of the attractive features of λ-calculus, such as being a rewriting system, especially one that is confluent or typable.

The SF-calculus achieves this by replacing the operator K of SK-calculus with a *factorisation operator* F that is able to test the internal structure of terms that are, in some sense, head normal, by factoring them. The operator F takes three arguments. If O is an operator and M and N are combinators then $FOMN$ reduces to M, so that FF represents the traditional K. If PQ is a *compound* then $F(PQ)MN$ reduces to NPQ, so that N can manipulate the components separately. It cannot be stressed too much that not every application is a compound. Rather, compounds are applications that can never be reduced at their head. That is, they are given by all partially applied operators such as SM and SMN and F and FMP for any terms M, N and P. Non-examples include fully applied operators such as $SMNP$ and $FMNP$, and terms headed by a variable, such as xM. The latter is excluded since substitution may create a fully applied operator. Using factorisation, closed normal forms can be completely analysed into their constituent operators, whose equality can be tested by extensional means. Details of the calculus can be found in the original paper [34].

Theorem 8 *The normal model of computability for SF-calculus is equivalent to the recursive model.*

Proof That Gödelisation and Church encoding are both simulations follows from the work of Church [16] and Kleene [41], so that it is enough to show that both re-codings are computable. It is easy to see that the recoding of numbers to numbers is recursive. In the other direction, the recoding of SF-combinators can be described by a pattern-matching function that acts on the combinators in normal form. Such pattern-matching functions are represented by SF-combinators because SF-calculus is *structure complete* [34]. Note that this proof does not apply for SK-calculus as this is merely *combinatorially complete* [67]. Also, since the recodings are invertible, this is an equivalence of partial combinatory algebras, in the sense of Longley [45]. □

To the extent that λ-calculus is equivalent to SK-calculus, this makes SF-calculus a superior foundation for higher-order programming languages.

9 Conclusions

The Church-Turing Thesis has been a confusion since it was first named, but not defined, by Kleene in 1952. The numerical results of Church and Turing support numerical versions of their eponymous theses, in which sets of numerical functions are co-extensive. Further, there are separate theses for symbolic

DRAFT

computation, involving simulations of one model of computability in another. Kleene confused these two settings, with a little encouragement from Church.

Once the role of simulations is made explicit, it is easier to see that mutual simulation yields an equivalence of models only if both re-codings are computable, each in its respective model. This requirement exposes the limitations of λ-calculus, since Gödelisation is not λ-definable, even for closed λ-terms in normal form.

These limitations are, in some sense, well known within the λ-calculus community, in that λ-calculus cannot define equality, even of closed normal forms. Indeed, those working with categorical models of computability, or analysing programs defined as λ-terms, are acutely aware of these limitations. However, the community as a whole is not keen to advertise any of this, proclaiming instead that λ conquers all. Students who ask the wrong questions may be told "Beware the Turing tarpit!" [54] or "Don't look under the lambda!" which closes off discussion without clarifying anything.

The limitations of λ-calculus are essential to its nature, since λ-terms cannot directly query the internal structure of their arguments; the expressive power of λ-calculus is extensional. This does not matter for numerical computations since the internal structure of a natural number is determined by the zero-test and the predecessor function, both of which are recursive. However, this approach cannot be generalised to query internal structure in richer settings.

Rather, intensional computation requires a fresh outlook. The simplest illustration of this is the SF-calculus whose factorisation operator F is able to uniformly decompose normal forms to their constituent operators. Since SF-calculus also has all of the expressive power of SK-calculus, its normal model of computability is equivalent to the Turing model or the recursive function model.

The implications of this for programming language design are profound. The **bondi** programming language has already shown how the usual database queries can be made polymorphic, and that object-orientation can be defined in terms of pattern-matching. Now the factorisation operator paves the way for program analysis to be conducted in the source language, so that growing a language can become easier than ever.

In short, confusion in the Church-Turing Thesis has obscured the fundamental limitations of λ-calculus as a foundation for programming languages. It is time to wind up Landin's research program, and pursue the development of intensional calculi and programming languages.

Acknowledgments We thank Jacques Carette, Robin Cockett, John Crossley, Nachum Dershowitz, Thomas Given-Wilson, Neil Jones, Jens Palsberg, Reuben Rowe, Peter Selinger, and Eric Torreborre for comments on drafts of this paper.

References

[1] Martin Abadi and Luca Cardelli. *A Theory of Objects*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1st edition, 1996.

[2] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, December 2000.

DRAFT

[3] Sergei Artemov and Eduardo Bonelli. The intensional lambda calculus. In Sergei N. Artemov and Anil Nerode, editors, *Logical Foundations of Computer Science*, volume 4514 of *Lecture Notes in Computer Science*, pages 12–25. Springer Berlin Heidelberg, 2007.

[4] Henk P. Barendregt. *The Lambda Calculus — Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, revised edition, 1984.

[5] Henk P. Barendregt and Giulio Manzonetto. Turing's contributions to lambda calculus. In B. Cooper and J. van Leeuwen, editors, *Alan Turing — His Work and Impact*, pages 139–143. Elsevier, Boston, 2013.

[6] Henk P. Barendregt, Giulio Manzonetto, and Marinus J. Plasmeijer. The imperative and functional programming paradigm. In B. Cooper and J. van Leeuwen, editors, *Alan Turing — His Work and Impact*, pages 121–126. Elsevier, Boston, 2013.

[7] Ulrich Berger and Helmut Schwichtenberg. An inverse of the evaluation functional for typed λ-calculus. In Gilles Kahn, editor, *IEEE Symposium on Logic in Computer Science (LICS)*, pages 203–211. IEEE Computer Society Press, July 1991.

[8] Inge Bethke. *Notes on Partial Combinatory Algebras*. PhD thesis, Universiteit van Amsterdam, 1988.

[9] Udi Boker and Nachum Dershowitz. Comparing computational power. *Logic Journal of the IGPL*, 4(5):633–647, 2006.

[10] bondi programming language. `www-staff.it.uts.edu.au/~cbj/bondi`, 2014.

[11] Ana Bove, Peter Dybjer, and Ulf Norell. A brief overview of Agda — a functional language with dependent types. In *Proceedings of the 22nd International Conference on Theorem Proving in Higher Order Logics*, TPHOLs '09, pages 73–78, Berlin, Heidelberg, 2009. Springer-Verlag.

[12] L.E. Jan Brouwer. *Brouwer's Cambridge Lectures on Intuitionism*. Cambridge University Press, Cambridge ; New York, 1981. Edited by Dirk van Dalen.

[13] Jacques Carette and Aaron Stump. Towards typing for small-step direct reflection. In *Proceedings of the ACM SIGPLAN 2012 Workshop on Partial Evaluation and Program Manipulation*, PEPM '12, pages 93–96, New York, NY, USA, 2012. ACM.

[14] Carlo Cellucci. Proof theory and theory of meaning. In MariaLuisa Dalla Chiara, editor, *Italian Studies in the Philosophy of Science*, volume 47 of *Boston Studies in the Philosophy of Science*, pages 13–29. Springer Netherlands, Amsterdam, 1981.

[15] Donald D. Chamberlin and Raymond F. Boyce. Sequel: A structured English query language. In *Proceedings of the 1974 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*, SIGFIDET '74, pages 249–264, New York, NY, USA, 1974. ACM.

DRAFT

[16] Alonzo Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):pp. 345–363, 1936.

[17] Alonzo Church. *The calculi of lambda conversion*. Number 6 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, USA, 1985.

[18] J. Robin B. Cockett and Pieter J. W. Hofstra. Introduction to Turing categories. *Annals of Pure and Applied Logic*, 156(2-3):183–209, 2008.

[19] J. Robin B. Cockett and Pieter J. W. Hofstra. Categorical simulations. *Journal of Pure and Applied Algebra*, 214(10):1835–1853, 2010.

[20] Nicolas de Bruijn. A survey of the project AUTOMATH. In Jonathon P. Seldin and John R. Hindley, editors, *To H.B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism*, pages 579–606. Academic Press, Amsterdam, 1980.

[21] Solomon Feferman. Arithmetization of metamathematics in a general setting. *Fundamenta Mathematicae*, 49(1):35–92, 1960.

[22] Solomon Feferman. Turing's thesis. *Notices of the American Mathematical Society*, 53:1200–1206, 2006.

[23] Matthias Felleisen. On the expressive power of programming languages. In Neil Jones, editor, *ESOP '90*, volume 432 of *Lecture Notes in Computer Science*, pages 134–151. Springer Berlin, Heidelberg, 1990.

[24] Gottlob Frege. On sense and reference. In Peter Geach and Max Black, editors, *Translations from the Philosophical Writings of Gottlob Frege*, pages 56–78. Basil Blackwell, Oxford, 1960. Originally published in 1892 as "Über Sinn und Bedeutung." Translated by Max Black.

[25] Robin Gandy. Church's thesis and principles for mechanisms. In K. Jon Barwise, H. J. Keisler, and K. Kunen, editors, *The Kleene Symposium*, pages 123–148. North-Holland, Amsterdam, 1980.

[26] Thomas Given-Wilson. An intensional concurrent faithful encoding of Turing machines. In Ivan Lanese, Alberto Lluch Lafuente, Ana Sokolova, and Hugo Torres Vieira, editors, Proceedings 7th Interaction and Concurrency Experience, Berlin, Germany, 6th June 2014, volume 132 of *Electronic Proceedings in Theoretical Computer Science*, pages 19–35. Open Publishing Association, 2014.

[27] Thomas Given-Wilson, Daniele Gorla, and Barry Jay. A Concurrent Pattern Calculus. *Logical Methods in Computer Science*, 10(3), 2014.

[28] J. A. Goguen, J. W. Thatcher, E. G. Wagner, and J. B. Wright. Initial algebra semantics and continuous algebras. *Journal of the ACM*, 24(1):68–95, January 1977.

[29] Mayer Goldberg. Gödelization in the lambda calculus. *Information Processing Letters*, 75(1):13–16, 2000.

DRAFT

[30] Robert Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, New York, NY, USA, 2012.

[31] Jerry R. Hobbs and Stanley J. Rosenschein. Making computational sense of Montague's intensional logic. *Artificial Intelligence*, 9:287–306, 1978.

[32] J. Martin Hyland. The effective topos. In A.S. Troelstra and D. van Dalen, editors, *The L. E. J. Brouwer Centenary Symposium Proceedings of the Conference held in Noordwijkerhout*, volume 110 of *Studies in Logic and the Foundations of Mathematics*, pages 165–216. Elsevier, Amsterdam, 1982.

[33] Barry Jay. *Pattern Calculus: Computing with Functions and Data Structures*. Springer, 2009.

[34] Barry Jay and Thomas Given-Wilson. A combinatory account of internal structure. *J. Symb. Log.*, 76(3):807–826, 2011.

[35] Barry Jay and Delia Kesner. First-class patterns. *Journal of Functional Programming*, 19(2):34 pages, 2009.

[36] Barry Jay and Jens Palsberg. Typed self-interpretation by pattern matching. *ACM SIGPLAN Notices*, 46(9):247–258, September 2011.

[37] Barry Jay and Jose Vergara. Growing a language in pattern calculus. In *Theoretical Aspects of Software Engineering (TASE), 2013 International Symposium on*, pages 233–240, July 2013.

[38] C.B. Jay. The pattern calculus. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 26(6):911–937, November 2004.

[39] Neil D. Jones. *Computability and Complexity: From a Programming Perspective*. MIT Press, Cambridge, MA, USA, 1997.

[40] Jerrold J. Katz. The end of Milleanism: Multiple bearers, improper names, and compositional meaning. *The Journal of Philosophy*, 98(3):137–166, 2000.

[41] Stephen C. Kleene. λ-definability and recursiveness. *Duke Mathematical Journal*, 2:340–353, 1936.

[42] Steven C. Kleene. *Introduction to Metamathematics*. Bibliotheca Mathematica. North-Holland Publishing Company, Amsterdam, 1952. Copublisher: Wolters–Noordhoff; 8th revised ed.1980.

[43] Georg Kreisel. Some reasons for generalizing recursion theory. In *Logic Colloquium '69*, pages 139–198, Amsterdam, 1969. North–Holland Publishing Co.

[44] Peter J. Landin. The Next 700 Programming Languages. *Communications of the ACM*, 9:157–166, 1966.

[45] John Longley. Computability structures, simulations and realizability. *Mathematical Structures in Computer Science*, 24(02):e240201, 2014.

DRAFT

[46] Per Martin-Löf. *Intuitionistic Type Theory.* Bibliopolis, Napoli, 1984.

[47] John McCarthy. History of Lisp. *SIGPLAN Notices*, 13(8):217–223, August 1978.

[48] Robin Milner, Mads Tofte, and David MacQueen. *The Definition of Standard ML.* MIT Press, Cambridge, MA, USA, 1997.

[49] John C. Mitchell. On abstraction and the expressive power of programming languages. *Science of Computer Programming*, 21(2):141–163, 1993.

[50] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 10(3):470–502, July 1988.

[51] Torben Æ. Mogensen. Gödelization in the untyped lambda-calculus. In Olivier Danvy, editor, *Proceedings of the 1999 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation, San Antonio, Texas, USA, January 22-23, 1999. Technical report BRICS-NS-99-1*, pages 19–24. University of Aarhus, 1999.

[52] Richard Montague. Towards a general theory of computability. *Synthese*, 12(4):429–438, 1960.

[53] Reinhard Muskens. Intensional models for the theory of types. *Journal of Symbolic Logic*, 72(1):98–118, 2007.

[54] Alan J. Perlis. Special feature: Epigrams on programming. *SIGPLAN Notices*, 17(9):7–13, September 1982.

[55] Gordon D. Plotkin. The origins of structural operational semantics. *Journal of Logic and Algebraic Programming*, 60:3–15, 2004.

[56] Andrew Polonsky. Axiomatizing the quote. In Marc Bezem, editor, *Computer Science Logic (CSL'11) — 25th International Workshop/20th Annual Conference of the EACSL*, volume 12 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 458–469, Dagstuhl, Germany, 2011. Schloss Dagstuhl—Leibniz-Zentrum füer Informatik.

[57] Emil L. Post. Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics*, 65(2):pp. 197–215, 1943.

[58] John C. Reynolds. The essence of ALGOL. In Peter W. O'Hearn and Robert D. Tennent, editors, *ALGOL-like Languages, Volume 1*, pages 67–88. Birkhauser Boston Inc., Cambridge, MA, USA, 1997.

[59] Hartley Jr. Rogers. *Theory of Recursive Functions and Effective Computability.* MIT Press, Cambridge, MA, USA, 1987.

[60] Dana S. Scott. Data types as lattices. *SIAM Journal on Computing*, 5:522–587, 1976.

[61] Dana S Scott. λ-calculus: Then & now. In *ACM Turing Centenary Celebration*, page 9. ACM, 2012.

DRAFT

[62] Robert I. Soare. Computability and recursion. *The Bulletin of Symbolic Logic*, 2(3):pp. 284–321, 1996.

[63] Robert I. Soare. Chapter 1 The history and concept of computability. In Edward R. Griffor, editor, *Handbook of Computability Theory*, volume 140 of *Studies in Logic and the Foundations of Mathematics*, pages 3–36. Elsevier, 1999.

[64] G.L. Steele. Growing a language. *Higher-Order and Symbolic Computation*, 12(3):221–236, 1999.

[65] G. J. Sussman and G. L. Steele. SCHEME: An interpreter for extended lambda calculus. AI Memo 349, MIT Artificial Intelligence Laboratory, December 1975.

[66] Carolyn Talcott. *The essence of Rum — A Theory of Intensional and Extensional Aspects of Lisp Computation*. PhD thesis, Stanford University, August 1985.

[67] Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge U.K, 2003.

[68] Pavel Tichỳ. Intension in terms of Turing machines. *Studia Logica*, 24(1):7–21, 1969.

[69] Alan M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Procedings of the London Mathematical Society*, 42(2):230–265, 1936.

DRAFT

Contents

1 Introduction **1**

2 Church's Thesis **2**

3 Models of computability **4**

4 Comparison of models **5**

5 Turing's Thesis **7**

6 The Church-Turing Thesis **8**

7 Programming language design **9**

8 Intensional computation **11**

9 Conclusions **12**

DRAFT